arXiv:1602.04260v2 [cs.IT] 23 Jun 2016

Adaptivity provably helps: information-theoretic limits on l_0 cost of non-adaptive sensing

Sanghamitra Dutta and Pulkit Grover

Abstract

The advantages of adaptivity and feedback are of immense interest in signal processing and communication with many positive and negative results. Although it is established that adaptivity does not offer substantial reductions in minimax mean square error for a fixed number of measurements, existing results have shown several advantages of adaptivity in complexity of reconstruction, accuracy of support detection, and gain in signal-to-noise ratio, under constraints on sensing energy. Sensing energy has often been measured in terms of the Frobenius Norm of the sensing matrix. This paper uses a different metric that we call the l_0 cost of a sensing matrix– to quantify the complexity of sensing. Thus sparse sensing matrices have a lower cost. We derive information-theoretic lower bounds on the l_0 cost that hold for any non-adaptive sensing strategy. We establish that any non-adaptive sensing strategy must incur an l_0 cost of a $\Theta\left(N\log_2(N)\right)$ to reconstruct an N-dimensional, one–sparse signal when the number of measurements are limited to $\Theta\left(\log_2(N)\right)$. In comparison, bisection-type adaptive strategies only require an l_0 cost of at most $\mathcal{O}(N)$ for equal order of measurements. The problem has an interesting interpretation as a sphere packing problem in a multidimensional space, such that all the sphere centres have minimum non-zero co-ordinates. We also discuss the variation in l_0 cost as the number of measurements increase from $\Theta\left(\log_2(N)\right)$ to $\Theta\left(N\right)$.

Index Terms

Adaptive Compressed Sensing, Binary Search, Sparse Reconstruction, Matrix Norms, Sphere Packing

I. Introduction

THE problem of compressed sensing and sparse reconstruction [1] [2] has drawn tremendous interest over the past decade due to its potential applications in sub-Nyquist sampling, imaging, biomedical signal processing, astronomy and geophysics. Consider the problem of reconstructing an N dimensional sparse signal vector $\boldsymbol{x}$ from its noisy M dimensional compressed measurement vector $\boldsymbol{y}$:

$$\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x} + \boldsymbol{z} \tag{1}$$

Here $\boldsymbol{A}$ is the $M \times N$ dimensional sensing matrix with $M \ll N$ and $\boldsymbol{z}$ is the noise vector. We assume $\boldsymbol{z} \sim \mathcal{N}(0, \boldsymbol{I}_{M\times M})$. The signal $\boldsymbol{x}$ is said to be K–sparse (Usually $K \ll N$) if $||\boldsymbol{x}||_0 \leq K$. A detailed review of various sparse reconstruction algorithms proposed over the past decade can be found in [3].

It appears intuitive that during the acquisition of compressive measurements of an unknown signal $\boldsymbol{x}$, "choosing an adaptive or sequential strategy that cleverly selects the next rows of sensing matrix based on what has been previously observed" [4] might help in easily reconstructing the unknown signal as compared to a non-adaptive strategy. In a non-adaptive strategy, the entire measurement vector $\boldsymbol{y}$ is acquired in one-shot and the rows of sensing matrix $\boldsymbol{A}$ are not chosen based on any prior or acquired information. In this paper, we choose the l_0 cost (similar to "l_0 norm"), *i.e.*, number of non-zero entries of the sensing matrix as our measure of sensing complexity for comparison between adaptive [5] [6] and non-adaptive strategies.

In this paper, we derive a general information-theoretic lower bound on the l_0 cost of reconstruction of a one–sparse signal that hold for any non-adaptive sensing matrix and demonstrate that adaptivity can achieve a lower l_0 cost as compared to any non-adaptive strategy, by a factor that diverges to infinity for large N. It is well established that both adaptive [5] and non-adaptive [7] [8] sensing strategies require at least $\Theta\left(\log_2(N)\right)$ measurements [9] [10]for the reconstruction of a one–sparse signal. We show that in the reconstruction of a one–sparse signal, any adaptive strategy would require an l_0 cost of $\Theta\left(N\log_2(N)\right)$ as compared to bisection-type [5] adaptive strategies that require at most an $\mathcal{O}(N)$ l_0 cost when the number of measurements are limited to $\Theta(\log_2(N))$. Our problem translates into a novel sphere packing problem within an M dimensional sphere with an additional constraint of minimizing the l_0 cost of the sphere centres. We also discuss scenarios when the number of measurements varies from $\Theta(\log_2(N))$ to $\Theta(\log_2(N))$. Lastly, we provide outlines on extension of our results to K–sparse signals.

A sparser sensing matrix would offer significant advantages in terms of implementation, storage and reconstruction as discussed in [11]. However, this line of work is mostly focused on the design of non-adaptive, sparse matrices and efficient reconstruction algorithms, rather than the derivation of fundamental limits. Adding on to the benefits of sparse matrices in [11], one can view the process of acquiring a measurement, *i.e.*, computing the dot product of the signal with a row of the sensing matrix as a "filtering-type" operation. The magnitude of the entries in the sensing matrix then corresponds to amplification/scaling during sample acquisition, while a zero value implies that no filter-tap (multiplier) is turned on. Then the

Sanghamitra Dutta and Pulkit Grover are with Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA 15213, USA. Emails:- sanghamd, pgrover@andrew.cmu.edu.

This paper has been accepted at the IEEE International Symposium on Information Theory, 2016.

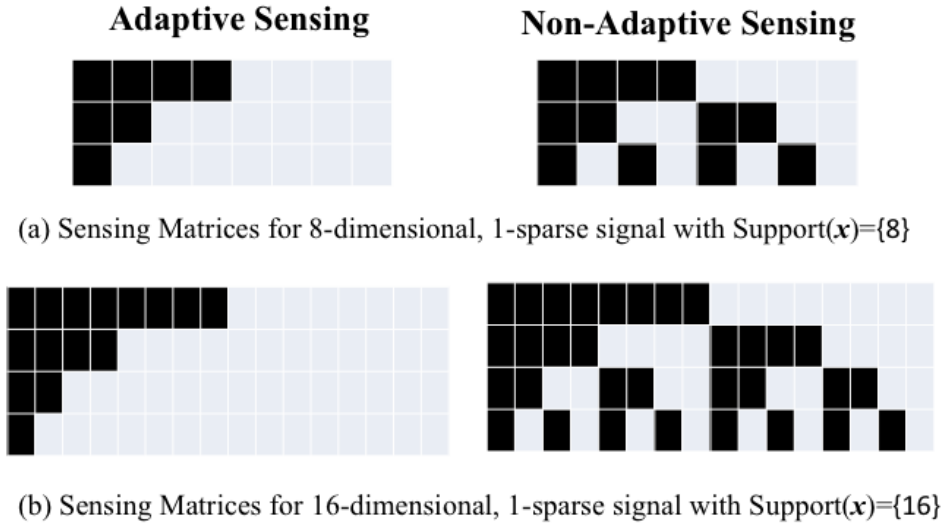
Adaptive Sensing
Non-Adaptive Sensing
(a) Sensing Matrices for 8-dimensional, 1-sparse signal with Support(x)={8}
(b) Sensing Matrices for 16-dimensional, 1-sparse signal with Support(x)={16}

Fig. 1. The figure provides intuition on why adaptive sensing could have fundamentally sparser matrices than nonadaptive sensing. The matrices on the left use an adaptive bisection strategy to narrow down the search space in $\Theta(\log_2 N)$ steps, requiring $\Theta(\log_2 N)$ measurements. Interestingly, the construction on the right can simulate the bisection strategy nonadaptively. However, the cost is increased density of the matrices.

l_0 cost would actually correspond to the number of multipliers required during compressive acquisition and thus have a direct implication on the sensing cost. A higher number of multipliers would also add more noise to the system and also increase the acquisition complexity. Similarly, in group testing/pooling type applications, the l_0 cost would correspond to the number of mixing operations required.

While not in the context of sparsity of sensing matrices, there is a body of literature that addresses the question of whether and when adaptivity is really helpful. Arias-Castro, Candes and Davenport [4] derive a minimax (worst case) lower bound on the mean square error for any sensing strategy (adaptive/non-adaptive). While they do not show any explicit advantage of adaptivity as in this paper, the authors argue that the potential benefits in Mean Square Error offered by adaptivity "in the worst case" are within $\mathcal{O}(\log(N))$ over non-adaptive, random strategies. In this paper, we show an explicit advantage to adaptivity as compared to *any* non-adaptive strategy and also highlight the costs of using random matrices in terms of l_0 cost. On the other hand, Malloy and Nowak [5] point out various advantages to adaptivity that include reducing reconstruction complexity (as does feedback in noisy channels) and also in providing gain in signal-to-noise ratio under limitations of sensing energy. They show that "bisection-type" adaptive strategies can recover the support of the unknown K–sparse signal with minimum amplitude $\Omega\left(\sqrt{\frac{N}{B}\log(K)}\right)$ as compared to traditional non-adaptive strategies that require an amplitude of $\Omega\left(\sqrt{\frac{N}{B}\log(N)}\right)$ [12] [13] for the same number of measurements, *i.e.*, $\Theta(K\log_2(N))$. Here $\sqrt{B}$ is a constraint on the Frobenius norm of the sensing matrix. In a different scenario where $M > N$, Haupt, Castro and Nowak [14] establish that adaptivity in sensing can recover signals with "vanishingly small" amplitudes as compared to random non-adaptive strategies, under sensing energy constraints. Some other related references that show advantages of adaptivity in specific scenarios can be found in [15], [16] and [17].

Why would one think adaptivity can make the sensing matrix sparser? Fig.1 illustrates a scenario of "bisection-type" adaptive sensing in comparison to non-adaptive sensing as the sensing matrices get sparser. It provides an intuition that while the number of measurements remain the same (at least in order sense), even with the sparsest non-adaptive sensing matrices, adaptivity always incurs a lower l_0 cost than non-adaptive sensing. References in favour of adaptivity [5] have cited existing information-theoretic bounds [13] [12] on SNR and number of measurements for non-adaptive compressive sensing that are mostly based on certain deterministic matrices or random, Gaussian matrices. However, Gaussian matrices are dense with an l_0 cost of NM. Adaptivity can allocate non-zero entries in the sensing matrix selectively, like a binary search. The intuition is most clear for one–sparse signals, and hence those are discussed first.

II. Notation and Problem Formulation

Let the l_0 cost of a sensing matrix $\boldsymbol{A} = [A_1 A_2 \dots A_N]$ in problem (1) be defined as:

$$||\boldsymbol{A}||_{0,0} = \sum_{i=1}^{N} ||A_i||_0 = \sum_{i=1}^{N}\sum_{j=1}^{M} \mathbb{I}(A(j,i) \neq 0) \tag{2}$$

Here A_i denotes the i-th column of $\boldsymbol{A}$ and $\mathbb{I}(.)$ denotes the indicator function. We say a function $F(N) = \Omega(g(N))$ if there exist a positive constant C_0 such that, $F(N) \geq C_0 g(N)$ as $N \to \infty$. It implies that asymptotically, $g(N)$ is not dominated by $F(N)$. Similarly, $F(N) = \Theta(g(N))$ if there exist positive constants C_1, C_2 such that, $C_1 g(N) \leq F(N) \leq C_2 g(N)$. It implies that neither $F(N)$ nor $g(N)$ are asymptotically dominated by each other. Lastly, $F(N) = \mathcal{O}(g(N))$ if $F(N) \leq C_3 g(N)$ for some positive constant C_3.

We derive lower bounds on $||\boldsymbol{A}||_{0,0}$ for non-adaptive sensing that asymptotically dominate upper bounds for adaptive sensing. In Theorem 1 we consider the case of reconstruction of a one sparse signal $\boldsymbol{x}$ using a *binary* sensing matrix $\boldsymbol{A}$ in the absence of noise. In Theorem 2, we extend our analysis to real-valued sensing matrices for reconstructing a one sparse signal $\boldsymbol{x}$ with a target probability of error in the presence of additive white Gaussian noise. Note that we have assumed M to vary as a logarithmic function of N, when N is large.

III. Main Results

Our first theorem establishes an information-theoretic lower bound on the l_0 cost of non-adaptive sensing for *binary* sensing matrices in the absence of noise when the number of measurements are limited to $\Theta\left(\log_2(N)\right)$.

Theorem 1: For the reconstruction of an arbitrary one–sparse signal with amplitude μ using *binary* sensing matrices with number of measurements $M = T\log_2(N)$ where T is a constant greater or equal to 1, any non-adaptive sensing strategy incurs an l_0 cost of $\Theta\left(N\log_2(N)\right)$.

Proof Sketch: Consider the recovery of the support of a one–sparse vector $\boldsymbol{x}$, from $\boldsymbol{y}$ and $\boldsymbol{A}$ when $M = \Theta\left(\log_2(N)\right)$. Then, any non-adaptive strategy would at least require the columns $\{A_i\}$ of $\boldsymbol{A}$ to be distinct. Notice

$$\boldsymbol{A}\boldsymbol{x} = [A_1, \cdot\ \cdot, A_N]\begin{bmatrix} x_1 \\ \cdot \\ \cdot \\ x_N \end{bmatrix} = \sum_{i=1}^{N} A_i x_i$$

If two columns (say i_1 and i_2) of $\boldsymbol{A}$ are equal, then two one–sparse $\boldsymbol{x}$ vectors with respective supports $\{i_1\}$ and $\{i_2\}$ become indistinguishable. Among all *binary* sensing matrices $\boldsymbol{A}$ with N unique columns, we consider one with minimum l_0 cost and derive lower bounds on its l_0 cost. Such a *binary* sensing matrix $\boldsymbol{A}$ with minimum l_0 cost can be easily designed by choosing the columns in ascending order of l_0 cost, *i.e.*, first all possible columns with l_0 cost 1, then all possible columns with l_0 cost 2 and so on until we get N unique columns. We show that $||\boldsymbol{A}||_{0,0}$ is always $\Theta(N\log_2(N))$.

Note that for $M < \log_2(N)$, it is not possible to construct binary sensing matrices with unique non-zero columns. So we assume $M = T\log_2(N)$ where $T \geq 1$. Our next theorem establishes a lower bound on the l_0 cost of any non-adaptive sensing matrix in the presence of noise, when $||A_i||_2$ is bounded. For non-adaptive strategies, limitations of sensing energy [5] usually enforce upper bounds on $||A_i||_2$, thus justifying the assumptions of *Theorem 2*.

Theorem 2: For the reconstruction of an arbitrary one–sparse signal with amplitude μ in the presence of Gaussian Noise of unit variance, using real-valued sensing matrices, under the constraints that $M = T\log_2(N)$ where T is a constant greater or equal to 1 and $||A_i||_2$ is bounded by τ, any non-adaptive sensing strategy cannot achieve a target probability of error ϵ without incurring an l_0 cost of $\Theta\left(N\log_2(N)\right)$, provided that the allowed bound on $||A_i||_2$ *i.e.*, $\tau = \Theta(Q^{-1}(\epsilon)/\mu)$.

Proof Sketch: First consider the case of one–sparse signal $\boldsymbol{x}$ with non-zero value equal to μ. Think of each column of $\boldsymbol{A}$ as a point in the M dimensional space $\mathbb{R}^M$. The minimum probability of error in distinguishing two points in $\mathbb{R}^M$, say A_i and A_j (that correspond to two columns of $\boldsymbol{A}$) scaled by μ is always bounded by $P_e^{i,j} \geq Q\left(\frac{\mu||A_i - A_j||_2}{2}\right)$ where $||\cdot||_2$ denotes the l_2 norm and $Q(\cdot)$ is the Gaussian tail function defined as $Q(a) = \frac{1}{\sqrt{2\pi}}\int_a^\infty e^{-b^2/2}db$. For any non-adaptive strategy to attain a target probability of error ϵ, we must have

$$||A_i - A_j||_2 \geq \frac{2Q^{-1}(\epsilon)}{\mu}\ \forall\ i, j \in \{1, 2, \ldots, N\},\ i \neq j \tag{3}$$

Let $d = \frac{2Q^{-1}(\epsilon)}{\mu}$. Under $||A_i||_2 \leq \tau$ for some constant τ, the problem translates to finding N unique points in a sphere in $\mathbb{R}^M$ of radius τ, that are each separated by at least d. As before, among all such matrices, (*i.e.*, a collection of N points in $\mathbb{R}^M$), we consider a matrix $\boldsymbol{A}$ with minimum l_0 cost and attempt to find a lower bound for it.

Corollary 1: Bisection-type adaptive strategies only require an l_0 cost of $\mathcal{O}(N)$ to reconstruct one–sparse signals when the number of measurements vary from $\Theta(\log_2(N))$ to $\Theta(N)$and are thus asymptotically dominated by non-adaptive strategies.

In the next Section, we provide some preliminary lemmas and then we use them prove our theorems.

IV. Proofs of Theorems 1 and 2

A. Preliminary Results

Lemma 1: Let r be an integer such that $1 \leq r < M/2$. Also, let $p = \frac{r}{M} < 1/2$.

$$\sum_{l=1}^{r}\binom{M}{l} \leq \frac{\binom{M}{r}}{1-2p} \tag{4}$$

Moreover, if there exists a constant λ such that $p \leq \lambda < 1/2$, then the partial summation is of the same order as the last term.

$$\binom{M}{r} \leq \sum_{l=1}^{r}\binom{M}{l} \leq \frac{\binom{M}{r}}{1-2\lambda} \tag{5}$$

Proof: First observe that

$$\frac{\binom{M}{r-1}}{\binom{M}{r}} = \frac{r}{M-r+1} < \frac{2r}{M} = 2p < 1 \tag{6}$$

Similarly, for $t = 1, 2, \cdots, r-2$,

$$\frac{\binom{M}{r-t-1}}{\binom{M}{r-t}} = \frac{r-t}{M-r+t+1} < \frac{2r}{M} = 2p < 1 \tag{7}$$

Thus,

$$\begin{aligned}\sum_{l=1}^{r}\binom{M}{l} &\leq \binom{M}{r}\left(1+(2p)+(2p)^2+\cdots+(2p)^{(r-1)}\right)\\ &\leq \binom{M}{r}\left(1+(2p)+(2p)^2+\cdots\right) = \frac{\binom{M}{r}}{1-2p}\end{aligned} \tag{8}$$

Note that, if $p \leq \lambda < 1/2$,(5) follows. ■

Lemma 2: [18, Lemma 4.7.1] The binary entropy function is defined as $H(p) = -p\log_2(p) - (1-p)log_2(1-p)$ where $0 \leq p \leq 1$. Let $p = \frac{r}{M} \leq 1$.

$$\frac{2^{MH(p)}}{\sqrt{8Mp(1-p)}} \leq \binom{M}{r} \leq \frac{2^{MH(p)}}{\sqrt{2\pi Mp(1-p)}} \tag{9}$$

Lemma 3: [19, Lemma 2] For an integer $d > 1$, and $p < 1/2$

$$H(p) \leq \frac{2d}{ln(2)}(p)^{1-\frac{1}{d}} \tag{10}$$

B. Binary Sensing Matrices

Proof of Theorem 1: Recall that we are required to bound $||\boldsymbol{A}||_{0,0}$ where the columns of $\boldsymbol{A}$ are chosen in ascending order of l_0 cost. Let r_0 be the integer such that

$$\sum_{l=1}^{r_0}\binom{M}{l} < N \leq \sum_{l=1}^{r_0+1}\binom{M}{l} \tag{11}$$

For $l = 1, 2, \cdots, r_0$, there can be at most $\binom{M}{l}$ columns with l ones, (*i.e.*, l_0 cost l) and the remaining $\left(N - \sum_{l=1}^{r_0}\binom{M}{l}\right)$ columns have l_0 cost (r_0+1).

If $M \leq \lfloor \log_2(N) \rfloor$, it is not possible to construct *binary* matrices with N unique columns. If $M = \lfloor \log_2(N) \rfloor + 1$, then $r_0 \leq M/2$ and $||\boldsymbol{A}||_{0,0} \approx 0.5N\log_2(N)$. As M increases beyond $\lfloor \log_2(N) \rfloor + 1$ for fixed N, r_0 decreases, and *so does* $||\boldsymbol{A}||_{0,0}$. Thus, it is sufficient to prove the lower bound on $||\boldsymbol{A}||_{0,0}$ for $M = T\log_2(N)$ where T is a constant strictly greater than 1, as $||\boldsymbol{A}||_{0,0}$ for $T = 1$ will only be higher. For $T > 1$, we always have $r_0 < M/2$ (required for *Lemma 1*).

Let $N = c_M \sum_{l=1}^{r_0}\binom{M}{l}$ for some c_M. Let $p_0 = \frac{r_0}{M}$. First, we show that if $M = T\log_2(N)$ ($T > 1$), then there exists a constant λ such that $p_0 \leq \lambda < 1/2$. Note that,

$$\frac{2^{MH(p_0)}}{\sqrt{8Mp_0(1-p_0)}} \overset{Lemma2}{\leq} \sum_{l=1}^{r_0}\binom{M}{l} < N \tag{12}$$

Taking logarithm of both sides,

$$H(p_0) \leq \frac{\log_2(N)}{M} + \frac{\log_2(\sqrt{8r_0})}{M} < \frac{1}{T} + \frac{\log_2(\sqrt{8M})}{M} \tag{13}$$

Note that, $\frac{\log_2(\sqrt{8M})}{M}$ is decreasing and tends to 0 for large M. For any constant λ' such that $(1/T) < \lambda' < 1$, there always exists an $M = M(\lambda')$, such that $\frac{\log_2(\sqrt{8M})}{M} < \lambda' - (1/T)$ for all $M > M(\lambda')$. Or, there exists a constant $\lambda = H^{-1}(\lambda')$ such that, $p_0 \leq \lambda < H^{-1}(1) = 1/2$. This is required to apply *Lemma 1*.

We now provide a lower bound on the l_0 cost of $\boldsymbol{A}$.

$$\begin{aligned}
||\boldsymbol{A}||_{0,0} &= \sum_{l=1}^{r_0} \binom{M}{l}\, l + \left(N - \sum_{l=1}^{r_0} \binom{M}{l}\right)\, (r_0+1) \\
&\geq r_0 \binom{M}{r_0} + \sum_{l=1}^{r_0} \binom{M}{l} (c_M - 1)(r_0+1) \\
&\geq r_0\ c_M \binom{M}{r_0} \overset{Lemma1}{\geq} r_0\ c_M \sum_{l=1}^{r_0} \binom{M}{l} (1-2\lambda) \\
&= (1-2\lambda)\ r_0\ N
\end{aligned} \tag{14}$$

Now, note that $r_0 < M$. Therefor either $r_0 = \Theta(M)$ or $r_0 = \boldsymbol{o}(M)$. We show here that $r_0 = \Theta(M)$. Look into the range of c_M.

$$1 < c_M \leq \frac{\sum_{l=1}^{r_0+1} \binom{M}{l}}{\sum_{l=1}^{r_0} \binom{M}{l}} < \frac{M+1}{r_0+1} < \frac{M}{2} \tag{15}$$

Using *Lemma 1* and *2*,

$$N = c_M \sum_{l=1}^{r_0} \binom{M}{l} \leq \frac{c_M 2^{MH(p_0)}}{(1-2\lambda)\sqrt{\pi r_0}} \tag{16}$$

Taking logarithm of both sides,

$$\begin{aligned}
\frac{\log_2(N)}{M} &\leq \frac{\log_2(c_M/((1-2\lambda)\sqrt{\pi r_0})}{M} + H(p_0) \\
&\overset{(15)}{\leq} \frac{\log_2(M/(1-2\lambda))}{M} + H(p_0)
\end{aligned} \tag{17}$$

Recall that $M = T\log_2(N)$ for some constant $T > 1$. For large M, $\frac{\log_2(M/(1-2\lambda))}{M} \to 0$. For any positive constant $\eta' < (1/T)$, there always exist an $M = M(\eta')$ such that $\frac{\log_2(M/(1-2\lambda))}{M} < \eta'$ for all $M > M(\eta')$.
Or, $H(p_0) \geq (1/T) - \eta' > 0$ for large M, which implies that, $r_0 \geq H^{-1}((1/T) - \eta')M = \eta M$ for some constant η. Using this result in (14), we finally obtain the lower bound,

$$||\boldsymbol{A}||_{0,0} \geq (1-2\lambda)\eta N M = CN\log_2(N) \tag{18}$$

Here $C = (1-2\lambda)\eta T$ is a constant. Needless to say, the l_0 cost $||\boldsymbol{A}||_{0,0} \leq NM = TN\log_2(N)$ which is the total number of elements in an $M \times N$ matrix. Thus the sensing matrix is always dense,*i.e.*, the l_0 cost is of same order as the total number of elements. Thus $||\boldsymbol{A}||_{0,0} = \Theta(N\log_2(N))$ ■

C. Real Valued Sensing Matrix in Noise

Proof of Theorem 2: The performance of non-adaptive reconstruction strategies relies on how well one can distinguish any two columns of sensing matrix $\boldsymbol{A}$, in the presence of noise. As already mentioned, we first consider the one–sparse case. We need to find N unique points in a sphere in $\mathbb{R}^M$ dimensional space of radius τ that are each separated by at least $d = \frac{2Q^{-1}(\epsilon)}{\mu}$. Assume we have u_1 points of l_0 cost 1, u_2 points of l_0 cost 2 and so on. Thus, $N = u_1 + u_2 + \cdots + u_{r_0'}$ and $||\boldsymbol{A}||_{0,0} = \sum_{l=1}^{r'} u_l\, l$.

Now consider the maximum possible number of points with l_0 cost equal to 1, and separated by distance d in an M dimensional sphere. Clearly there are $\binom{M}{1}$ possible choices of the non-zero dimension, and at most $\frac{2\tau}{d}$ points can be placed along each dimension. Thus $u_1 \leq \binom{M}{1}\frac{2\tau}{d}$.

Similarly, consider the maximum possible number of points with l_0 cost equal to 2, and separated by distance d. There are $\binom{M}{2}$ possible choices of 2 non-zero dimensions among M. For any such combination of 2 dimensions, the problem is equivalent to fitting maximum number of $2D$ spheres of radius $d/2$ within a $2D$ sphere of radius τ. We would always have less than $\frac{\pi\tau^2}{\pi(d/2)^2}$ such points (or equivalently spheres) along any such combination of 2 dimensions. Thus $u_2 \leq \binom{M}{2}\left(\frac{2\tau}{d}\right)^2$.

For any $l \in \{1, 2, \cdots, M\}$, the maximum number of points with l_0 cost equal to l and separated by distance d within the M dimensional sphere of radius τ is always less than $\binom{M}{l}(\frac{2\tau}{d})^l$ as shown in Fig.2. This follows since there are $\binom{M}{l}$ possible choices of the l non-zero dimensions and for each such choice of l non-zero dimensions, the number of l dimensional spheres of radius $d/2$ that can be placed in a big l dimensional sphere of radius τ is given by the ratio of the volumes of the two l dimensional spheres. Thus $u_l \leq \binom{M}{l}(\frac{2\tau}{d})^l$.

Choose $v_l = \binom{M}{l}\left(\frac{2\tau}{d}\right)^l\ \forall\ l \in 1, 2..M$. Let r_0 be the integer such that $\sum_{l=1}^{r_0} v_l < N \leq \sum_{l=1}^{r_0+1} v_l$. Thus, for $\kappa > 0$,
$v_1 + v_2 + ..v_{r_0} + \kappa = N = u_1 + u_2 + \cdots + u_{r_0'}$

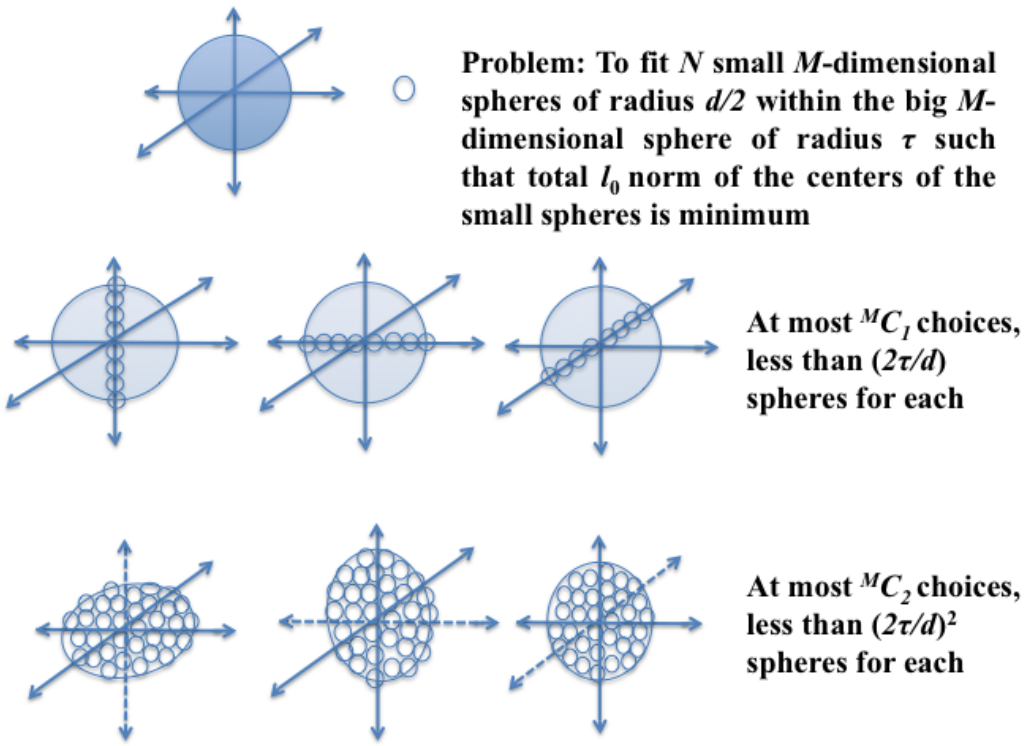
Problem: To fit N small M-dimensional spheres of radius d/2 within the big M-dimensional sphere of radius τ such that total l_0 norm of the centers of the small spheres is minimum
At most $^{M}C_1$ choices, less than (2τ/d) spheres for each
At most $^{M}C_2$ choices, less than (2τ/d)^2 spheres for each

Fig. 2. A New Sphere Packing Problem - The figure illustrates the problem of finding N unique columns of the sensing matrix $\boldsymbol{A}$ with minimum l_0 cost. The problem has an interesting interpretation as a sphere packing problem. For simplicity, we choose $M = 3$ in this figure and then show how we can keep placing spheres of radius $d/2$ such that the total l_0 cost of the centres of the spheres (columns of $\boldsymbol{A}$) is minimum.

First we prove that $r_0 < r_0'$. Assume $r_0 \geq r_0'$. Then,

$$\begin{aligned} N = u_1 + u_2 + .. + u_{r_0'} &\leq v_1 + v_2 + .. + v_{r_0'} \\ &\leq v_1 + v_2 .. + v_{r_0} < N \end{aligned} \tag{19}$$

This is a contradiction. Thus, the l_0 cost can be lower bounded as shown below, using the fact that the summation puts more weights on the lower cost terms.

$$||\boldsymbol{A}||_{0,0} = \sum_{l=1}^{r_0'} u_l \; l \geq \sum_{l=1}^{r_0} v_l l + \kappa(r_0 + 1) \tag{20}$$

We follow a similar analysis to find a lower bound for $\sum_{l=1}^{r_0} v_l \; l + \kappa(r_0 + 1)$. Let $c_M \sum_{l=1}^{r_0} v_l = N$ and $p_0 = r_0/M$.

$$H(p_0) + p_0 \log_2 \left(\frac{2\tau}{d}\right) \leq \frac{\log_2(N)}{M} + \frac{\log_2(\sqrt{8r_0})}{M} \rightarrow \frac{1}{T} \tag{21}$$

Recall that, $\left(\frac{2\tau}{d}\right)$ is $\Theta(1)$. Thus, for large M, there exists a constant $\lambda < 1/2$ such that, $p_0 \leq \lambda < H^{-1}(1) = 1/2$. We now provide a lower bound on the l_0 cost of $\boldsymbol{A}$.

$$\begin{aligned} ||\boldsymbol{A}||_{0,0} &\geq \sum_{l=1}^{r_0} v_l l + \sum_{l=1}^{r_0} v_l (c_M - 1) \; (r_0 + 1) \\ &\geq r_0 \; c_M \; v_{r_0} \overset{\textit{Lemma 1}}{\geq} r_0 N (1 - 2\lambda) \end{aligned} \tag{22}$$

Lastly, we show that $r_0 = \Theta(M)$. As before, since it is already known that $r_0 < M$, either $r_0 = \Theta(M)$ or $r_0 = \boldsymbol{o}(M)$. Look into the range of c_M.

$$1 < c_M \leq \frac{\sum_{l=1}^{r_0+1} v_l}{\sum_{l=1}^{r_0} v_l} < 1 + \frac{(M - r_0)(\frac{2\tau}{d})}{r_0 + 1} < M \frac{2\tau}{d} \tag{23}$$

Like the previous proof,

$$N = c_M \sum_{l=1}^{r_0} \binom{M}{l} \left(\frac{2\tau}{d}\right)^l \leq \frac{c_M 2^{MH(p_0)}}{(1-2\lambda)\sqrt{\pi r_0}} \left(\frac{2\tau}{d}\right)^{r_0} \tag{24}$$

Taking logarithm and using (23)

$$\frac{\log_2(N)}{M} \leq \frac{\log_2\left(\frac{2\tau M}{d(1-2\lambda)}\right)}{M} + H(p_0) + p_0 \log_2\left(\frac{2\tau}{d}\right) \tag{25}$$

Now, if $p_0 = r_0/M \to 0$ for large M, then all the terms in R.H.S. of (25) would tend to 0. But since $M = T\log_2(N)$ for some constant $T \geq 1$, the L.H.S. is always a non-zero constant $1/T$. Thus we arrive at a contradiction. Thus, we have $r_0 \geq \eta M$ for some constant $\eta = \Theta(1)$. Using this result in (22), we obtain,

$$||\boldsymbol{A}||_{0,0} \geq (1-2\lambda)\eta N M \geq CN\log_2(N) \tag{26}$$

As $M = \Theta(\log_2(N))$, we have $||\boldsymbol{A}||_{0,0} = \Theta(N\log_2(N))$. ■

D. Comparison with Adaptive Sensing

Proof of Corollary: Consider the l_0 cost of the bisection-type adaptive sensing strategies as described in [6]. As in Fig.1, first the total length $\{1, 2, \cdots, N\}$ is divided into two equal blocks, and any one measurement is acquired. Based on the received measurement value, the active block (set of possible non-zero locations) is narrowed down. That blocks is again divided into two parts and any one part is sensed. For such a sensing strategy,

$$||\boldsymbol{A}||_{0,0} \leq \left(\frac{N}{2} + \frac{N}{4} + \cdots\right) \leq N \tag{27}$$

Thus, adaptive sensing matrices require an l_0 cost of at most $\mathcal{O}(N)$ for signal reconstruction and are thus sparse. Clearly, this upper bound is asymptotically dominated by $N\log_2(N)$.

Note that, since we have already allowed the values of the non-zero elements of $\boldsymbol{x}$ to be sufficiently high, it is guaranteed (from [5]) to succeed in the presence of noise.

E. Discussion on higher orders of M

We consider the case of reconstruction of one–sparse signals using *binary* sensing matrices, when M is between $\log_2(N)$ and N. From (13), we note that $p_0 \to 0$ if $\log_2(N)/M \to 0$ for large M and N. We can always find a constant λ arbitrarily small, such that $p_0 \leq \lambda < 1/2$. The binomial summation in *Lemma 1* is practically equal to the largest term. From (14)

$$||\boldsymbol{A}||_{0,0} \geq c_M r_0 \binom{M}{r_0} \geq r_0 N(1-2\lambda) \approx r_0 N \tag{28}$$

In general, note that when $M = \Theta(\log_2(N))$, then $r_0 = \Theta(M)$ and $||\boldsymbol{A}||_{0,0} = \Theta(N\log_2(N))$. As M increases, the lower bound on l_0 cost decreases. When M is of higher order, say $M = N$, then $r_0 = 1$ and we have a reduced l_0 cost N using simply an identity sensing matrix. From (17),

$$\frac{\log_2(N/c_M)}{M} \leq H(r_0/M) \overset{\textit{Lemma 3}}{\leq} \frac{2d}{ln(2)} \left(\frac{r_0}{M}\right)^{1-\frac{1}{d}} \tag{29}$$

$$\left(\frac{ln2\ \log_2(N/c_M)}{2d}\right)^{d/(d-1)} M^{-1/(d-1)} \leq r_0 \tag{30}$$

This leads to the following bound.

$$||\boldsymbol{A}||_{0,0} \geq C \max_d \frac{N}{M^{1/(d-1)}} \left(\frac{ln2\log_2(N/c_M)}{2d}\right)^{d/d-1} \tag{31}$$

Similar bounds may be derived for real-valued matrices.

V. Conclusions

This paper clearly shows an advantage of adaptivity in terms of l_0 cost of sensing. It establishes a lower bound on the l_0 cost that holds for any non-adaptive strategy, whether it is random or not. K–sparse signals, we can argue that if a matrix does not have unique columns and fails to reconstruct one–sparse signals, it naturally fails to reconstruct a K–sparse signal. If a matrix has two similar columns say $\{i_1\}$ and $\{i_2\}$, then two exactly K sized supports differing only at $\{i_1\}$ and $\{i_2\}$ become indistinguishable. Thus, the lower bound on the l_0 cost of $\boldsymbol{A}$ cannot be less than that in the case of one–sparse signals. Thus, $||\boldsymbol{A}||_{0,0} = \Theta\left(NK\log_2(N)\right)$ when $M = \Theta(K\log_2(N))$. For small values of K, this still asymptotically dominates $\mathcal{O}(KN)$ which is the l_0 cost for bisection-type adaptive strategies.

ACKNOWLEDGMENT

This work was supported in part by Systems on Nanoscale Information fabriCs (SONIC), one of the six SRC STARnet Centers, sponsored by MARCO and DARPA. We also acknowledge the support of NSF-CCF-1350314 (NSF Career Award). S Dutta was also supported by a Dean's Fellowship from Carnegie Mellon University.

We also thank Praveen Venkatesh, Haewon Jeong, Yaoqing Yang and Aarti Singh for their suggestions on the paper.

REFERENCES

[1] E. Candès and M. Wakin, "An introduction to compressive sampling," *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 21–30, 2008.
[2] D. Donoho, "Compressed sensing," *IEEE Transactions on Information Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.
[3] J. Tropp and S. Wright, "Computational methods for sparse solution of linear inverse problems," *Proceedings of the IEEE*, vol. 98, no. 6, pp. 948–958, 2010.
[4] E. Arias-Castro, E. Candes, and M. Davenport, "On the fundamental limits of adaptive sensing," *IEEE Transactions on Information Theory*, vol. 59, no. 1, pp. 472–481, 2013.
[5] M. L. Malloy and R. D. Nowak, "Near-optimal adaptive compressed sensing," *IEEE Transactions on Information Theory*, vol. 60, no. 7, pp. 4001–4012, 2014.
[6] M. Davenport and E. Arias-Castro, "Compressive binary search," in *Proceedings IEEE International Symposium on Information Theory*, 2012, pp. 1827–1831.
[7] R. Baraniuk, M. Davenport, R. DeVore, and M. Wakin, "A simple proof of the restricted isometry property for random matrices," *Constructive Approximation*, vol. 28, no. 3, pp. 253–263, 2008.
[8] K. Do Ba, P. Indyk, E. Price, and D. P. Woodruff, "Lower bounds for sparse recovery." in *SODA*, vol. 10. SIAM, 2010, pp. 1190–1197.
[9] P. Grover, "Fundamental limits on power consumption for lossless signal reconstruction," in *IEEE Information Theory Workshop (ITW)*, 2012, pp. 527–531.
[10] C. Aksoylar and V. Saligrama, "Information-theoretic bounds for adaptive sparse recovery," in *Proceedings IEEE International Symposium on Information Theory*, 2014, pp. 1311–1315.
[11] A. Gilbert and P. Indyk, "Sparse recovery using sparse matrices," *Proceedings of the IEEE*, vol. 6, no. 98, pp. 937–947, 2010.
[12] W. Wang, M. J. Wainwright, and K. Ramchandran, "Information-theoretic limits on sparse signal recovery: Dense versus sparse measurement matrices," *IEEE Transactions on Information Theory*, vol. 56, no. 6, pp. 2967–2979, 2010.
[13] S. Aeron, V. Saligrama, and M. Zhao, "Information theoretic bounds for compressed sensing," *IEEE Transactions on Information Theory*, vol. 56, no. 10, pp. 5111–5130, 2010.
[14] J. Haupt, R. M. Castro, and R. Nowak, "Distilled sensing: Adaptive sampling for sparse detection and estimation," *IEEE Transactions on Information Theory*, vol. 57, no. 9, pp. 6222–6235, 2011.
[15] M. Malloy and R. Nowak, "On the limits of sequential testing in high dimensions," in *Proceedings Asilomar Conference on Signals, Systems and Computers*. IEEE, 2011, pp. 1245–1249.
[16] R. M. Castro, J. Haupt, R. Nowak, and G. M. Raz, "Finding needles in noisy haystacks," in *Proceedings IEEE International Conference on Acoustics, Speech and Signal Processing*, 2008, pp. 5133–5136.
[17] G. Braun, S. Pokutta, and Y. Xie, "Info-greedy sequential adaptive compressed sensing," in *Proceedings Allerton Conference on Communication, Control, and Computing*. IEEE, 2014, pp. 858–865.
[18] R. B. Ash, *Information Theory*. New York: Dover Publications, 1990.
[19] A. Sahai and P. Grover, "The price of certainty:" waterslide curves" and the gap to capacity," *arXiv preprint arXiv:0801.0352*, 2008.